



Benchmark Procurement Guidelines

for

Government PC Buyers

March, 2005 Edition

© 2002-2004 Advanced Micro Devices, Inc. All rights reserved.

The contents of this document are provided in connection with Advanced Micro Devices, Inc. (“AMD”) products. AMD makes no representations or warranties with respect to the accuracy or completeness of the contents of this publication and reserves the right to make changes to specifications and product descriptions at any time without notice. No license, whether express, implied, arising by estoppel, or otherwise, to any intellectual property rights are granted by this publication. Except as set forth in AMD’s Standard Terms and Conditions of Sale, AMD assumes no liability whatsoever, and disclaims any express or implied warranty, relating to its products including, but not limited to, the implied warranty of merchantability, fitness for a particular purpose, or infringement of any intellectual property right.

AMD’s products are not designed, intended, authorized or warranted for use as components in systems intended for surgical implant into the body, or in other applications intended to support or sustain life, or in any other application in which the failure of AMD’s product could create a situation where personal injury, death, or severe property or environmental damage may occur. AMD reserves the right to discontinue or make changes to its products at any time without notice.

Trademarks

AMD, the AMD Arrow logo, AMD Athlon, AMD Opteron, AMD Sempron, AMD Turion, and combinations thereof, are trademarks of Advanced Micro Devices, Inc.

HyperTransport is a licensed trademark of the HyperTransport Technology Consortium.

Windows is a registered trademark of Microsoft Corporation.

SYSmark, Business Winstone and Content Creation Winston are registered trademarks of Ziff Davis, Inc.

SPEC is a registered trademark of Standard Performance Evaluation Corporation.

Other product names used in this publication are for identification purposes only and may be trademarks of their respective companies.

THE USE OF BENCHMARKS IN TECHNICAL SPECIFICATIONS FOR THE PROCUREMENT OF PERSONAL COMPUTERS, NOTEBOOKS AND SERVERS

In many countries the technical specifications in public tenders for procuring computers (personal computers hereafter referred to as PCs, notebooks, and servers) often include particular brand or product names to describe microprocessors. These specifications do not comply with most current government procurement policies. Furthermore, such specifications discourage price competition and place government purchases at risk by basing purchasing decisions on brand names instead of accurate measures of computer performance.

- **The use of brand or product names in government procurement specifications is illegal:** Many governments, including the European Union and the United States, have laws and regulations in place that generally prohibit the use of brand or product names in government tenders. A similar prohibition is provided in the World Trade Organisation (WTO) Government Procurement Agreement (GPA). Additionally, the courts, antitrust authorities and contracting authorities in many countries around the world have ruled that such specifications are prohibited.
- **The use of brand or product names or even certain technical specifications is not an accurate measure of computer performance:** Discrimination does not benefit the customer because brand or product names do not describe product performance and quality. Further, certain technical specifications such as clock rate as measured in MHz or GHz, or cache size, or front side bus speed, etc. tend to be brand and/or architecture specific and do not accurately describe product performance or quality.
- **The use of brand or product names inhibits competition and cost savings:** Specifications with a specific brand or product name discourage the cost savings that accompanies competition. Neutral specifications encourage vendors to offer innovative products at competitive prices.

The use of Benchmarks to Describe Computer Performance

To ensure that government agencies enjoy both the benefits of fair and open competition and comply with procurement regulations, the use of accurate performance measures for computers and specifications that are brand-neutral, objective, and based on performance are necessary. The most effective means to accomplish this is for computer specifications to describe benchmark scores based on the rigorous testing methods developed by qualified independent third-party benchmarking firms or industry-standard consortiums that are widely recognized by the semiconductor computer and consumer electronics industries.

The advantage of using benchmarks is that they provide the potential buyer with a specific, easily comparable measure for the performance of a computer in an objective way. In addition, the use of benchmarks allows contracting authorities to draft simple yet accurate technical performance specifications for computers based on a defined benchmark score.

In general, there are two kinds of benchmarks:

1. Application-based benchmarks which measure the performance of computer systems by testing actual software applications in usage scenarios representative of how a computer is actually used;
2. Synthetic benchmarks which measure performance of computers by using non-application based tests that either try to estimate the performance of actual software applications or try to isolate a specific subsystem of a computer and measure its performance.

Although synthetic benchmarks are often easier and faster test to perform, they rarely correlate to how real software applications perform and rarely stand the test of time given the complexity and diversity of

computer systems and computer subsystems. Given this, AMD strongly recommends technical performance specifications should, whenever possible, be based on application-based benchmarks.

Benchmarking Computer Performance for Desktop PCs and Notebooks

Given the rapid rate of change in computer technology, the underlying architecture of the microprocessor, and application software, benchmarks are generally updated on an annual basis. Currently, BAPCo's SYSmark® 2004 is recommended for its accuracy, objectivity, simplicity, and widespread adoption. Alternatively, if for some reason SYSmark 2004 is not utilized, Business Winstone® 2004, Content Creation Winstone® 2004, or Worldbench 5.0 are recommended. Note that since desktop and notebook PCs are multi-purpose devices and application software can be used in so many ways, to improve the accuracy of a specification the geometric mean of two or more of the above benchmarks is further recommended.

BAPCo SYSmark® 2004: BAPCo is a non-profit consortium that develops and distributes a set of objective performance benchmarks for PCs and notebooks based on popular software applications and operating systems. BAPCo's current members include: AMD, ARCintuition, ATi Technologies, CNET, Computer Shopper, Dell, Hewlett-Packard, Intel, Microsoft, nVidia, Seagate, Toshiba, Transmeta, VNU, VNU Labs Ziff Davis Media, and ZDNet. SYSmark 2004 measures the performance of computers by running real applications through a series of scripted operations. Benchmark scoring is calculated on the basis of the time required to perform all the operations for the two usage scenarios described below.

- **SYSmark 2004 - Office Productivity** is BAPCo's latest version of the mainstream office productivity benchmark. This benchmark contains scientifically designed workloads that represent a range of activities that an office productivity worker may encounter.
- **SYSmark 2004 – Internet Content Creation** is BAPCo's latest version of the internet content creation benchmarks. This benchmark measures performance of the computer in relation to the typical applications used when creating, for instance, web-sites, or multimedia images and/or documents.

The aggregate SYSmark 2004 score is recommended, although the Office Productivity or the Internet Content Creation score may also be used. Additional information can be found at www.bapco.com.

VeriTest Business Winstone® 2004 and Multimedia Content Creation Winstone® 2004: VeriTest is a division of Lionbridge Technologies Inc., a private company which provides outsourced testing solutions. Both tests measure the performance of computers by running through a series of scripts on today's top-selling Windows® operating system-based applications. Benchmark scoring is calculated on the basis of the time required to performance all the operations for the two usage scenarios described below.

- **Business Winstone 2004** is Veritest's latest version of mainstream business application benchmarks.
- **Multimedia Content Creation Winstone 2004** is Veritest's latest the latest version multimedia content creation benchmarks

Additional information on Business Winstone 2004 can be found at www.veritest.com.

PC World Worldbench 5.0: PC World Communications, Inc., is a subsidiary of International Data Group, the world's leading technology media, research, and event company. WorldBench is an applications-based benchmark that automatically installs a series of programs on your computer and then runs various tests to see how those apps perform on your system. The individual times for each test are used to create a composite score that indicates your system's overall performance relative to other systems that have also run WorldBench. Additional information can be found at www.pcworld.com.

Benchmarking Computer Performance for Servers and Workstations

Because of the diverse and often targeted applications used on servers and workstations, it is important to utilize benchmarks that are designed for the specific applications that the contracting authority is most likely to use. Fortunately, a rich set of industry-standard consortiums and other development/auditing organizations exist to aid in the performance specification. Two leading organizations are:

Standard Performance Evaluation Corporation (SPEC®): SPEC is a non-profit corporation formed to establish, maintain and endorse a standardized set of relevant benchmarks that can be applied to the newest generation of high-performance computers. SPEC develops suites of benchmarks and also reviews and publishes submitted results from its member organizations and other benchmark licensees. For more information see www.spec.org.

Transaction Processing Performance Council (TPC): The TPC is a non-profit corporation founded to define transaction processing and database benchmarks and to disseminate objective, verifiable TPC performance data to the industry. For more information see www.tpc.org.

Given changes in computer technology, the underlying architecture of the microprocessor, and application software, benchmarks are regularly updated. The following benchmarks are currently recommended. More information can be found in the respective links.

- Database Servers
 - TPC-C (<http://www.tpc.org/tpcc/>)
 - TPC-H (<http://www.tpc.org/tpch/>)
 - SAP Standard Application Benchmarks (<http://www50.sap.com/benchmark/>)
- Messaging Servers
 - MMB3 (<http://www.microsoft.com/exchange/evaluation/performance/mmb3.asp>)
 - Lotus NotesBench (<http://www.notesbench.org/bench.nsf>)
- Web Servers
 - SPECweb99 (<http://www.spec.org/web99/>)
 - SPECweb99_SSL (<http://www.spec.org/web99ssl/>)
 - WebBench (<http://www.veritest.com/benchmarks/webbench/default.asp>)
- Java Client/Servers
 - jAppServer2004 (<http://www.spec.org/jAppServer2004/>)
 - JBB2000 (<http://www.spec.org/jbb2000/>)
- File Servers
 - NetBench (<http://www.veritest.com/benchmarks/netbench/default.asp>)
- Computational Servers and Workstations
 - SPEC CPU2000 (<http://www.spec.org/cpu2000/>)
 - HPC 2002 (<http://www.spec.org/hpc2002/>)
 - OMP 2001 (<http://www.spec.org/omp/>)
 - Fluent (<http://www.fluent.com/software/fluent/fl5bench/>)
- Graphics Workstations
 - SPECviewperf 8.0 (<http://www.spec.org/gpc/opc.static/vp8info.html>)
 - SPECapc for 3ds max 6 (<http://www.spec.org/gpc/apc.static/max6info.html>)
 - SPECapc for Maya 5 (<http://www.spec.org/gpc/apc.static/maya5info.html>)
 - SPECapc for pro/ENGINEER 2001 (<http://www.spec.org/gpc/apc.static/proe2001info.html>)
 - SPECapc for Solid Edge V14 (<http://www.spec.org/gpc/apc.static/se14info.html>)
 - SPECapc for SolidWorks 2003 (<http://www.spec.org/gpc/apc.static/sw2003.html>)

APPENDIX 1: Clock Rate Does Not Equal Performance

Technical specifications that describe the performance required by a microprocessor by using only one specific feature (for example, the clock rate) are discriminatory since one specific feature alone cannot correctly assess the performance required from the microprocessor and typically favors only a specific kind of microprocessor.

As an example, the specification of the **clock rate** of a microprocessor (in MHz or GHz) is not in itself indicative of the performance of the microprocessor. Put simply, the performance of the microprocessor is the product of the clock rate multiplied by the Instructions per clock (**IPC**). A microprocessor with a high IPC but low clock rate may therefore be a better product than a competing processor with a low IPC and a high clock rate.

Technical specifications including a combination of different features such as, for example, the clock rate, the front side bus speed and the level of cache, in an attempt to describe the performance required from a computer can equally be discriminatory since trade-offs from one feature to the other can be very difficult to calibrate. Consequently, specifications such as these usually favor a specific kind of microprocessor and discriminate against all others.

To be fair and open, technical specifications must be architecturally neutral.

APPENDIX II: Discrimination is Costly

Competition is the only way to achieve savings in procurement, as well as access a larger and more diverse product supply from a greater number of producers. That's why an increasing number of contracting authorities are identifying specific benchmarks as the best legal way to define the required PC performance in the public tenders; e.g. the European Commission is pushing all 25 member countries to declare a mandatory benchmark as the only legal measure for IT public procurement.

APPENDIX III: How to Use Industry Standard Benchmarks

With regard to the use of benchmarks, contracting authorities have a choice between several different methods.

One method consists of consulting one of the reputable lists that are published and constantly updated by benchmarks creators or publishers ("benchmark publications"). These benchmark publications often include descriptions of many different configurations of computers and/or servers and the corresponding benchmark score achieved by each configuration. The contracting authority should locate the benchmark score of its preferred configuration among the ones in the lists and use it as a reference for the performance required by the computer to be procured.

Alternatively, if the contracting authority does not find a suitable computer or server configuration in the benchmark publications, the contracting authority can ask one or more of the computer manufacturers ("OEMs") to provide him with the benchmark score of the desired configuration to be used as a reference in the tender to be issued.

If the contracting authority has no idea about the desired configuration, the authority can also check the benchmark publication or ask one of the OEMs for the benchmark score of its current machines and then require in the tender computers or servers that achieve a benchmark score exceeding by X% the score of the current machines.

Finally, a contracting authority can directly acquire one (or more) of the benchmarks available on the market and perform autonomously all the testing desired to find the benchmark score of the most suitable configuration.